

Compiling cooperative multitasking of ABS to Scala

Andri Saar, Keiko Nakata
Institute of Cybernetics at Tallinn University of Technology
`andri@cs.ioc.ee`, `keiko@cs.ioc.ee`

Abstract

In order to prove the correctness of compilation for our new Scala-based distributed backend for ABS, we will focus on one aspect of the ABS language, namely cooperative multitasking inside concurrent object groups. We create a simplified language that includes this feature and then define a compilation function that targets a language based on Scala that supports delimited continuations and prove the correctness of the compilation function.

1 Introduction

The ABS language [2], developed by the HATS project, is an Abstract Behavioral Specification language for concurrent object-based systems, designed to straddle the gap between design-oriented and implementation-oriented specification languages. It allows high-level specification of a system, by abstracting over implementation details while remaining executable and providing flexible concurrency and synchronization mechanisms by means of asynchronous method calls, separation of objects into concurrent groups and cooperative scheduling for objects inside one group. As the language serves as a foundation for several different analysis, it also has a formally defined semantics.

The concurrency model of ABS can be split in two: in one layer, we have local, synchronous and shared-memory communication between objects in one concurrent object group (COG) and on the second layer we have asynchronous message-based concurrency between different concurrent object groups reminiscent of Creol [3]. The behaviour of one COG is based on the cooperative multitasking of external method invocations and internal method activations, with concrete scheduling points where a different task may get scheduled. Between different COGs only asynchronous method calls may be used; different COGs have no shared object heap. The order of execution of asynchronous method calls is not specified. The result of an asynchronous call is a future; callers may decide at runtime when to synchronize with the reply from a call. Asynchronous calls may be seen as triggers that spawn new method activations (or tasks) within objects. Every object has a set of tasks that are to be executed (originating from method calls). Among these, at most one task of all the objects belonging to one COG is active; others are suspended and awaiting execution.

As ABS is designed to be an executable language, there are several backends available that are able to execute programs: the reference implementation is defined in Maude, and there exists a Java-based backend designed to help in developing and debugging ABS models. We are creating a new backend for ABS, using the Scala language, that would provide true distribution to ABS by allowing different concurrent object groups to be executed on different nodes in the network. As the concurrency models of ABS and Scala are quite different, we employ the experimental continuations plugin for Scala [4] to implement cooperative multitasking within one thread of execution.

In this article, we focus on the first step of providing a provably correct code generator for ABS. We look at what happens inside one concurrent object group and define a simplified language based on ABS that lacks objects but retains asynchronous calls and cooperative

scheduling between tasks. For this language we define a compilation function into a target language inspired by Scala that lacks concurrency but provides delimited continuations and closures. We then proceed to show that after compiling programs in the source language to a program in the target language preserves the operational behaviour.

2 Implementing the ABS concurrency model in Scala

As the concurrency model of ABS is notably different from the one provided by Scala, we need to provide a runtime support library for our compiler that would provide the necessary communication primitives to the generated code.

The asynchronous message-based communication between concurrent object groups is similar to actors. In practice, objects may pass three kinds of values between COGs: (a) object references, (b) future references and (c) immutable data structures from the functional sublanguage. Our implementation has to guarantee that all of these values are easily serializable so they could be passed around over the network, and make sure that the references are unique in the whole cluster and addressable from any node.

Inside one COG we have cooperative multitasking with the control being released only at certain *scheduling points*: in particular, `await e` statements. The statement `await e` suspends the execution of the current task until the guard expression e evaluates to true. This statement is a scheduling point: the scheduler may decide to switch tasks and proceed with some other task even if the guard immediately holds¹. In order to implement this efficiently we use continuations [1, 5]. Every time we reach a scheduling point, we capture the rest of the task in a continuation, store it together with its guard in the task set and give control back to the scheduler of the COG. The scheduler will then go through its list of tasks, evaluate the guards and picks a new task whose guard evaluates to true. If there are no such tasks, then none of the tasks in the COG are executable at that moment. In this case, the COG releases the execution thread until some external event causes some task to be ready for execution.

3 Compilation of simplified ABS

We have designed two languages, both endowed with small-step operational semantics, that are simpler than both ABS and Scala, yet exhibit the behaviour we are interested in. The source language models one COG in ABS by having asynchronous calls, `await` statement and cooperative multitasking. The target language includes primitive statements to manage the task set and supports continuations as the method to switch between tasks, which are used to implement a scheduler.

3.1 Source language

The source language is designed to be as minimalistic as possible while still retaining the interesting properties: thus, we have removed objects, COGs and futures from the language, resulting in a simple imperative programming language with asynchronous procedure calls and a way to wait for some condition to become true.

The statements S of the language are defined by:

$$S ::= x := e \mid \text{skip} \mid S_1; S_2 \mid \text{await } e \mid \text{if } e \text{ then } S_1 \text{ else } S_2 \mid \text{while } e \text{ do } S \mid f(\bar{e})$$

¹This is a simplification: in ABS, if the guard immediately holds, we are not allowed to switch tasks.

$$\begin{array}{c}
\frac{\text{env}_F \vdash n, \sigma \triangleright Ts' \rightarrow n', \sigma' \triangleright Ts''}{\text{env}_F \vdash n, \sigma \triangleright Ts \parallel Ts' \rightarrow n', \sigma' \triangleright Ts \parallel Ts''} \quad \frac{\llbracket e \rrbracket_{\rho', \sigma} = \mathbf{true}}{\text{env}_F \vdash n, \sigma \triangleright n \langle \rho \rangle \parallel n' \langle e, S, \rho' \rangle \rightarrow n', \sigma \triangleright n \langle \rho \rangle \parallel n' \langle S, \rho' \rangle} \\
\frac{}{\text{env}_F \vdash n, \sigma \triangleright n \langle \mathbf{await} \ e, \rho \rangle \rightarrow n, \sigma \triangleright n \langle e, \mathbf{skip}, \rho \rangle} \quad \frac{\text{env}_F \ f = (\bar{x}, S) \quad \bar{v} = \llbracket \bar{e} \rrbracket_{\rho, \sigma} \quad n' \text{ fresh}}{\text{env}_F \vdash n, \sigma \triangleright \langle f(\bar{e}), \rho \rangle \rightarrow n, \sigma \triangleright n \langle \mathbf{skip}, \rho \rangle \parallel n' \langle \mathbf{true}, S, [\bar{x} \mapsto \bar{v}] \rangle} \\
\frac{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1, \rho \rangle \rightarrow n \langle S'_1, \rho' \rangle}{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1; S_2, \rho \rangle \rightarrow n \langle S'_1; S_2, \rho' \rangle} \quad \frac{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1, \rho \rangle \rightarrow n \langle e, S'_1, \rho' \rangle}{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1; S_2, \rho \rangle \rightarrow n \langle e, S'_1; S_2, \rho' \rangle}
\end{array}$$

Figure 1: Structural operational semantics for the source language (excerpt)

Besides statements, we have pure expressions e which consist of usual arithmetic and boolean expressions. The notation \bar{e} (resp. \bar{x} and \bar{v}) denotes a possibly empty sequence of expressions (resp. variables and values).

A program P consists of a function environment env_F which maps function names to pairs of a formal argument list and a statement, and global variable declarations with their initial values. The entry point of the program will be the procedure named *main*.

The scheduler may only switch between tasks when evaluating an `await` e statement or when the currently running task terminates. Procedure calls $f(\bar{e})$ return immediately with no result and the original task will proceed. Communication between different tasks is achieved via global variables (which may be seen as analogues to object fields in ABS).

During runtime, the program is represented as a *configuration* c that consists of an active task identifier n , a global variable mapping σ and a set of tasks Ts . A task has an identifier and may be in one of the three forms: (a) a triple $\langle e, S, \rho \rangle$, representing a task that is awaiting to be scheduled, where e is the guard expression and ρ is the local variable mapping; (b) a pair $\langle S, \rho \rangle$, representing the currently active task and (c) a singleton $\langle \rho \rangle$, representing a terminated task.

$$\begin{array}{ll}
\textit{Configuration} \quad c & ::= n, \sigma \triangleright Ts \\
\textit{Task sets} \quad Ts & ::= n \langle e, S, \rho \rangle \mid n \langle S, \rho \rangle \mid n \langle \rho \rangle \mid Ts \parallel Ts
\end{array}$$

The order of tasks in the task set is irrelevant: the parallel operator \parallel is commutative and associative. Transition rules in the semantics are in the form $\text{env}_F \vdash c \rightarrow c'$. A subset of the rules is shown in Figure 1. We assume given an evaluation function $\llbracket e \rrbracket_{\rho, \sigma}$, which evaluates e with respect to the local and global variable mappings ρ and σ .

In the semantics, we are interested only in a subset of the whole task set and the first rule allows us to focus on the subset that we are interested in. If the currently active task is in the form $\langle S, \rho \rangle$, then we proceed to evaluate the statement S . Otherwise, it has terminated, i.e., it is in the form $\langle \rho \rangle$, or it has reached a scheduling point, i.e., in the form $\langle e, S, \rho \rangle$. Then we may pick another task from the task set whose guard evaluates to `true`. When a procedure f is invoked, we evaluate the argument expressions \bar{e} and put a new task $n' \langle \mathbf{true}, S, [\bar{x} \mapsto \bar{v}] \rangle$ to the task set, where S is the body of the procedure, $[\bar{x} \mapsto \bar{v}]$ maps formal argument names \bar{x} to their values \bar{v} and n' is a fresh identifier. The execution will proceed in the active task and the new task may only get scheduled once we reach a scheduling point.

3.2 Target language

The target language is modeled on Scala. Instead of containing explicit support for scheduling points, the target language supports delimited continuations via the `shift` and `reset` statements and the scheduler is explicitly specified when compiling programs from the source language to the target language. In order to not clutter the language with arrays, we also

3.3 Compilation function

The compilation function generates a program in the target language for every program in the source language. The interesting part of the definition is given by

$$\begin{aligned} \llbracket \text{fun } f(\bar{x}) \text{ is } S \rrbracket &= \text{fun } f(\bar{x}) \text{ is } \{\llbracket S \rrbracket; \text{invoke } sched\} \\ \llbracket \text{await } e \rrbracket &= \text{shift } k \{\text{addClosure } e \ k; \text{invoke } sched\} \\ \llbracket f(\bar{e}) \rrbracket &= \text{addFunc } f \ \bar{e} \end{aligned}$$

Every compiled program contains a prelude, which declares a global variable *sched* that will contain the scheduler continuation and a statement that adds the *main* function to the task set. The control is then passed on to the scheduler:

```
var sched : Closure = [skip, []];
addFunc main ();
reset {while isTask do shift k {sched = k; getTask()}}
```

The scheduler checks if there is any task ready to be executed (*isTask*); if there is, the rest of the computation (which means invoking the *while* loop again) is stored in the global variable *sched*. When the task terminates, the closure is invoked and the control passes back to the scheduler, which will proceed by trying to find another suitable task for execution. Similar pattern is followed with the *await e* statement: we store the rest of the execution in a continuation, add the continuation with its guard to the task set and pass the control back to the scheduler.

The following theorem states that the compilation function is correct: an one-step reduction in the source language is simulated by possibly multiple-step reduction in the target language after the compilation. To state the theorem formally, we extend the compilation function to configurations and define a similarity relation \preceq , which disregards the number of adjoining resets, between configurations in the target language.

Theorem 1 (Correctness of compilation). *For all configurations c_S, c'_S in the source language and c_T in the target language such that $c_S \rightarrow c'_S$ and $\llbracket c_S \rrbracket \preceq c_T$, there exists a configuration c'_T in the target language such that $c_T \mapsto^+ c'_T$ and $\llbracket c'_S \rrbracket \preceq c'_T$.*

4 Conclusion

We have designed a compilation function that compiles a simplified ABS-like language to a language based on Scala and showed that the resulting programs preserve the operational behaviour of the original program. This work can be used as a basis for extending it with other features currently omitted (such as futures, full objects and different COGs), gradually moving closer to ABS and its semantics as the source. The target semantics will also move closer to the actual implementation of the compiler and the Scala runtime. We also plan to formalise the proof with a proof assistant, such as Coq or Agda.

Acknowledgements

This research was supported by the EU FP7 project 231620 Highly Adaptable and Trustworthy Software (HATS) and the European Regional Development Fund (ERDF) through the Estonian Centre of Excellence in Computer Science (EXCS).

References

- [1] Olivier Danvy and Andrzej Filinski. Abstracting control. In *Proceedings of the 1990 ACM conference on LISP and functional programming*, LFP '90, pages 151–160, New York, NY, USA, 1990. ACM.
- [2] Einar Johnsen, Reiner Hähnle, Jan Schäfer, Rudi Scattee, and Martin Steffen. ABS: A core language for abstract behavioral specification. *Lecture Notes in Computer Science*. Springer, to appear.
- [3] Einar Johnsen and Olaf Owe. An asynchronous communication model for distributed concurrent objects. *Software and Systems Modeling*, 6:39–58, 2007. 10.1007/s10270-006-0011-2.
- [4] Tiark Rompf, Ingo Maier, and Martin Odersky. Implementing first-class polymorphic delimited continuations by a type-directed selective cps-transform. *SIGPLAN Not.*, 44:317–328, August 2009.
- [5] Mitchell Wand. Continuation-based multiprocessing. In *Proceedings of the 1980 ACM conference on LISP and functional programming*, LFP '80, pages 19–28, New York, NY, USA, 1980. ACM.
- [6] Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Inf. Comput.*, 115:38–94, November 1994.